

AMENDMENTS TO THE CLAIMS

1 – 93. (Cancelled).

94. (Currently Amended) A method for presenting materials corresponding to a navigation state, the method comprising:

receiving a user selection of an expression of attribute-value pairs;

producing a plurality of refinement options and a plurality of ancestors by, in each of a plurality of servers, processing the expression of attribute-value pairs to produce at least one refinement option and at least one ancestor;

combining the plurality of refinement options and plurality of ancestors to form combined refinement options, the combined refinement options including at least one refinement navigation state;

determining a navigation state associated with the expression of attribute-value pairs;

providing materials associated with the navigation state; and

providing the combined refinement options,

wherein the combining comprises:

taking a union of the plurality of refinement options,

determining a set of ancestors for each of the plurality of refinement options, from the plurality of ancestors produced in the plurality of servers, to form sets of ancestors,

computing an intersection of all of the sets of ancestors, and

computing the combined refinement options based on terms in the intersection of all sets of ancestors.

95. (Currently Amended) The method of claim 94, wherein the computing of the combined refinement options further comprising comprises identifying at least two related terms among the plurality of refinement options, and computing, for the at least two related terms, a least common ancestor of the related terms.

96. (Previously Presented) The method of claim 95, wherein the least common ancestor is defined by the partial order among the related terms.

97. (Previously Presented) The method of claim 95, wherein computing the least common ancestor of the related terms comprises storing all non-leaf terms on at least one server.

98. (Previously Presented) The method of claim 97, wherein the at least one server is a master server.

99. (Previously Presented) The method of claim 94, wherein the plurality of servers are slave servers, and wherein combining the plurality of refinement options includes storing a plurality of the attribute-value pairs on a master server.

100. (Cancelled).

101. (Previously Presented) The method of claim 94, wherein combining the plurality of refinement options is performed whether the plurality of refinement options are conjunctive, disjunctive, or negational.

102. (Previously Presented) The method of claim 94, wherein combining the plurality of refinement options comprises combining the plurality of refinement options as a disjoint union.

103. (Previously Presented) The method of claim 94, wherein the expression of attribute-value pairs is processed on a different partition of a collection of materials for different ones of the plurality of servers.

104. (Previously Presented) The method of claim 94, further comprising storing the attribute-value pairs in a graph data structure including nodes and edges between nodes, the nodes representing navigation states, the edges representing transitions.

105. (Previously Presented) The method of claim 94, further comprising partitioning the materials and other materials among the plurality of servers.

106. (Previously Presented) The method of claim 105, further comprising nesting the plurality of servers hierarchically.

107. (Previously Presented) The method of claim 106, wherein a root server of the plurality of servers acts as a master server and some of the plurality of servers act as slave servers, further comprising the master server distributing a request for a navigation state on to a plurality of slave servers, the slave servers computing navigation states for those requests and returning the results to the master server, and the master server combining the results from the slave servers to obtain a navigation state corresponding to the request.

108. (Previously Presented) The method of claim 94, further comprising storing the navigation state in a cache.

109. (Currently Amended) A computer-readable storage medium storing computer-executable instructions that, when executed by a computer, cause the computer to perform a method for presenting materials corresponding to a navigation state, the method comprising:

receiving a user selection of an expression of attribute-value pairs;

producing a plurality of refinement options and a plurality of ancestors by, in each of a plurality of servers, processing the expression of attribute-value pairs to produce at least one refinement option and at least one ancestor;

combining the plurality of refinement options and plurality of ancestors to form combined refinement options, the combined refinement options including at least one refinement navigation state;

determining a navigation state associated with the expression of attribute-value pairs;

providing materials associated with the navigation state; and

providing the combined refinement options,
wherein the combining comprises:
taking a union of the plurality of refinement options,
determining a set of ancestors for each of the plurality of refinement options, from
the plurality of ancestors produced in the plurality of servers, to form sets of ancestors,
computing an intersection of all of the sets of ancestors, and
computing the combined refinement options based on terms in the intersection of all
sets of ancestors.

110. (Currently Amended) The computer-readable storage medium of claim 109, wherein
the computing of the combined refinement options further ~~comprising~~ comprises identifying at least
two related terms among the plurality of refinement options, and computing, for the at least two
related terms, a least common ancestor of the related terms.

111. (Previously Presented) The computer-readable storage medium of claim 110,
wherein the least common ancestor is defined by the partial order among the related terms.

112. (Previously Presented) The computer-readable storage medium of claim 110,
wherein computing the least common ancestor of the related terms comprises storing all non-leaf
terms on at least one server.

113. (Previously Presented) The computer-readable storage medium of claim 112,
wherein the at least one server is a master server.

114. (Previously Presented) The computer-readable storage medium of claim 109,
wherein the plurality of servers are slave servers, and wherein combining the plurality of refinement
options includes storing a plurality of the attribute-value pairs on a master server.

115. (Cancelled).

116. (Previously Presented) The computer-readable storage medium of claim 109, wherein combining the plurality of refinement options is performed whether the plurality of refinement options are conjunctive, disjunctive, or negational.

117. (Previously Presented) The computer-readable storage medium of claim 109, wherein combining the plurality of refinement options comprises combining the plurality of refinement options as a disjoint union.

118. (Previously Presented) The computer-readable storage medium of claim 109, wherein the expression of attribute-value pairs is processed on a different partition of a collection of materials for different ones of the plurality of servers.

119. (Previously Presented) The computer-readable storage medium of claim 109, wherein the method further comprises storing the attribute-value pairs in a graph data structure including nodes and edges between nodes, the nodes representing navigation states, the edges representing transitions.

120. (Previously Presented) The computer-readable storage medium of claim 109, wherein the method further comprises partitioning the materials and other materials among the plurality of servers.

121. (Previously Presented) The computer-readable storage medium of claim 120, wherein the method further comprises nesting the plurality of servers hierarchically.

122. (Previously Presented) The computer-readable storage medium of claim 121, wherein a root server of the plurality of servers acts as a master server and some of the plurality of servers act as slave servers, the method further comprising the master server distributing a request for a navigation state on to a plurality of slave servers, the slave servers computing navigation states

for those requests and returning the results to the master server, and the master server combining the results from the slave servers to obtain a navigation state corresponding to the request.

123. (Previously Presented) The computer-readable storage medium of claim 109, wherein the method further comprises storing the navigation state in a cache.

124. (Currently Amended) A system for presenting materials corresponding to a navigation state, comprising:

at least one server that each:

receive a user selection of an expression of attribute-value pairs; and

produce at least one refinement option by processing the expression of attribute-value pairs;

at least one other server that:

receives a user selection of an expression of attribute-value pairs;

produces at least one refinement option by processing the expression of attribute-value pairs, wherein the at least one refinement option produced by the at least one server and the at least one refinement option produced by the at least one other server form a plurality of refinement options;

combines the plurality of refinement options to form combined refinement options, the combined refinement options including at least one refinement navigation state;

determines a navigation state associated with the expression of attribute-value pairs;

provides materials associated with the navigation state; and

provides the combined refinement options,

wherein when the at least one server combines the plurality of refinement options,
the at least one server:

takes a union of the plurality of refinement options,

determines a set of ancestors for each of the plurality of refinement options, from the plurality of ancestors produced in the plurality of servers, to form sets of ancestors,

computes an intersection of all of the sets of ancestors, and

computes the combined refinement options based on terms in the intersection of all sets of ancestors.

125. (Currently Amended) The system of claim 124, wherein ~~at least one of the at least one server and the at least one other server also~~ the computing of the combined refinement options comprises identifying ~~identifies~~ at least two related terms among the plurality of refinement options, and computing, for the at least two related terms, a least common ancestor of the related terms.

126. (Previously Presented) The system of claim 125, wherein the least common ancestor is defined by the partial order among the related terms.

127. (Previously Presented) The system of claim 125, wherein computing the least common ancestor of the related terms comprises storing all non-leaf terms the at least one of the at least one server and the at least one other server.

128. (Previously Presented) The system of claim 127, wherein the at least one of the at least one server and the at least one other server is a master server.

129. (Previously Presented) The system of claim 124, wherein the at least one server is a slave server, and wherein combining the plurality of refinement options includes storing a plurality of the attribute-value pairs on a master server.

130. (Cancelled).

131. (Previously Presented) The system of claim 124, wherein combining the plurality of refinement options is performed whether the plurality of refinement options are conjunctive, disjunctive, or negational.

132. (Previously Presented) The system of claim 124, wherein combining the plurality of refinement options comprises combining the plurality of refinement options as a disjoint union.

133. (Previously Presented) The system of claim 124, wherein the expression of attribute-value pairs is processed on a different partition of a collection of materials for different ones of the at least one server and the at least one other server.

134. (Previously Presented) The system of claim 124, wherein the attribute-value pairs are stored in a graph data structure including nodes and edges between nodes, the nodes representing navigation states, the edges representing transitions.

135. (Previously Presented) The system of claim 124, wherein the materials and other materials are partitioned among the at least one server and the at least one other server.

136. (Previously Presented) The system of claim 135, wherein the at least one server and the at least one other server are nested hierarchically.

137. (Previously Presented) The system of claim 136, wherein:
at least one of the at least one other server is a root server and acts as a master server;
the at least one server includes a plurality of servers, and the plurality of servers act as slave servers;

the master server distributes a request for a navigation state on to the slave servers;
the slave servers compute navigation states for those requests and return the results to the master server; and

the master server combines the results from the plurality of slave servers to obtain a navigation state corresponding to the request.

138. (Previously Presented) The system of claim 124, further comprising storing the navigation state in a cache.

139. (New) The method of claim 94, further comprising storing at least a portion of the plurality of ancestors produced in the plurality of servers in a cache.

140. (New) The method of claim 139, wherein at least one of the sets of ancestors is determined from the plurality of ancestors stored in the cache.

141. (New) The computer-readable storage medium of claim 109, further comprising storing at least a portion of the plurality of ancestors produced in the plurality of servers in a cache.

142. (New) The computer-readable storage medium of claim 141, wherein at least one of the sets of ancestors is determined from the plurality of ancestors stored in the cache.

143. (New) The system of claim 124, wherein the at least one server stores at least a portion of the plurality of ancestors produced in the plurality of servers in a cache.

144. (New) The system of claim 144, wherein at least one of the sets of ancestors is determined from the plurality of ancestors stored in the cache.